arXiv:0904.2906v1 [stat.ME] 19 Apr 2009

Sparse Bayesian Hierarchical Modeling of High-dimensional Clustering Problems

Heng Lian

Division of Mathematical Sciences

School of Physical and Mathematical Sciences

Singapore 637371

Singapore

Abstract

Clustering is one of the most widely used procedures in the analysis of microarray data, for example with the goal of discovering cancer subtypes based on observed heterogeneity of genetic marks between different tissues. It is well-known that in such high-dimensional settings, the existence of many noise variables can overwhelm the few signals embedded in the high-dimensional space. We propose a novel Bayesian approach based on Dirichlet process with a sparsity prior that simultaneous performs variable selection and clustering, and also discover variables that only distinguish a subset of the cluster components. Unlike previous Bayesian formulations, we use Dirichlet process (DP) for both clustering of samples as well as for regularizing the high-dimensional mean/variance structure. To solve the computational challenge brought by this double usage of DP, we propose to make use of a sequential sampling scheme embedded within Markov chain Monte Carlo (MCMC) updates to im-

prove the naive implementation of existing algorithms for DP mixture models. Our method is demonstrated on a simulation study and illustrated with the leukemia gene expression dataset.

keywords: Dirichlet process; Markov chain Monte Carlo; Sequential sampling; Sparsity prior;

1 Introduction

Clustering is one of the most widely used procedures in the analysis of microarray data. It has been used, for example, for cancer subtype discovery (Golub et al., 1999). Technological advances over the last decade on microarrays have made possible simultaneous investigation of thousands of genes that potentially characterize and distinguish previously known or unknown cancer subtypes. Although obviously not all the genes arrayed possess discriminative power for different cancer subtypes, if fewer genes are used, the procedure might fail to distinguish between some of the subtypes. Also, because of the cost of arraying the transcripts, this is a typical "large p, small n" problem that has attracted much attention recently.

Among many classes of clustering procedures, the model-based approach (Banfield and Raftery, 1993; Fraley and Raftery, 2002), assuming the data come from a mixture of distributions, has the advantage of permitting principled statistical inferences compared to other procedures based largely on heuristics, such as k-means. This is especially important in our case where inferences should be made on the selected variables as well as clustering structure.

Motivated by model interpretation as well as parsimony considerations, variable selection in clustering, mostly within the Bayesian framework, has been of increasing interest. Compared to variable selection in regression (Tibshirani, 1996; George and McCulloch, 1993, 1997), the clustering problem is much less studied.

Friedman and Meulman (2004) proposed one approach to select different subsets of variables and different associated weights for different clusters for non-model-based clustering. Liu et al. (2003) proposed to first reduce dimension by performing the principal component analysis on the covariates and then fitting a Bayesian mixture model to the top factors. Although the number of factors is automatically determined by the model, there are several disadvantages to this approach, including difficulty in interpretation in terms of the original attributes. Also, it can be argued that the top principal components do not necessarily have the most significant discriminative power for clustering and thus the procedure is suboptimal. Tadesse et al. (2005) adapted the stochastic search strategy of George and McCulloch (1993, 1997) originally proposed in the regression context and used reversible jump MCMC for inferences of cluster structures with simultaneous variable selection. This approach assumes that the same subsets of covariates discriminate all clusters. The model laid out in Kim et al. (2006) is based on the same philosophy but utilized an infinite mixture model via the Dirichlet process (DP) mixture. On the other hand, Hoff (2006) adopted a mean shift approach in which each cluster-specific mean deviates from the global base-line mean on one or more attributes that differs from cluster to cluster.

In this paper, we propose a Bayesian model for simultaneous clustering and variable selection via DP mixture as well. Our formulation is based on the mean shift model (Hoff, 2006). However, we use a novel hierarchical sparsity prior similar to that of Lucas et al. (2006) which can improve separation of significant signals from noise variables and thus can lead to reduced false discoveries. Also, we use a Dirichlet process shrinkage approach for both high-dimensional mean and variance that outperforms shrinkage using a non-DP prior, typically with normal distribution for mean and inverse-Gamma distribution for variance. Because of this double usage of Dirichlet process, both for sample clustering and for covariate shrinkage, the direct implementation of standard DP algorithms available in the literature becomes very

inefficient. We solve this problem by utilizing an embedded sequential sampling step as the proposal distribution in the Markov chain Monte Carlo (MCMC) iterations. In the next section, we formulate our model using the sparsity prior. Posterior computation via MCMC is discussed in detail in Section 3 where we also show how to use sequential sampling for efficient updating. For clarity in exposition, these two sections only consider shifts in means. Extension to shifts in both means and variances is briefly considered in Section 4. Section 5 includes a simulation study as well as an application to the leukemia gene expression data. We conclude the article with a brief discussion in Section 6.

2 Model Formulation

In this section as well as the next, we consider the case where the clusters differ from each other only in terms of their respective means for some of the attributes. In our model we start by expressing the samples $\mathbf{y}_i = (y_{i1}, \ldots, y_{ip}), i = 1, \ldots, n$, as

$$y_{ij} = \mu_j + \mu_{ij} + \sigma_j \epsilon_{ij},\ \epsilon_{ij} \overset{i.i.d.}{\sim} N(0,1).$$

In this formulation, μ_j and σ_j are attribute-specific mean and standard deviations shared by all samples. We put the following priors for them:

$$\mu_j \overset{i.i.d.}{\sim} DP(\alpha N(\mu_0, \sigma_0)),$$

$$\sigma_j^2 \overset{i.i.d}{\sim} DP(\beta Inv - Gamma(\alpha_0, \beta_0)),$$

where $DP(\alpha H)$ is the Dirichlet process with concentration parameter $\alpha > 0$ and base probability measure H. In this paper, we use the notation $\theta_i \overset{i.i.d.}{\sim} DP(\alpha H)$ as a short form for the more rigorous $\theta_i \overset{i.i.d.}{\sim} G, G \sim DP(\alpha H)$. This might be a misuse

but simplifies our notation since DP appears multiple times at different places within our model. When $\alpha \to \infty$, the first expression above reduces to $\mu_j \sim N(\mu_0, \sigma_0)$, for example. The use of Dirichlet process can be motivated from at least two point of views. First, it relaxes the normality assumption imposed on the components of the mean vector. Second, since the DP is a discrete measure, it provides a regularization mechanism by shrinking different parameters towards each other.

Since the attribute specific μ_j and σ_j are shared by all samples, the clustering structure can only derive from appropriate specification on μ_{ij}. As in Hoff (2006); Kim et al. (2006), the clustering of samples will be determined by an infinite mixture of distributions via Dirichlet process mixture. Denote $\boldsymbol{\mu}_i = (\mu_{i1}, \ldots, \mu_{ip})$. When it is intended that $\boldsymbol{\mu}_i$ is the mean for cluster c, i.e. sample i is assigned to cluster c, we also use $\boldsymbol{\mu}_c$ to denote the same mean vector. Although there might be some concern over misuse of notation, this can hardly cause any confusion in the context. The sample means are generated from an infinite mixture specified as the following:

$$\boldsymbol{\mu}_i = (\mu_{i1}, \ldots, \mu_{ip}) \overset{i.i.d.}{\sim} DP(\tau H),$$

where the base measure H on $\boldsymbol{\mu}_i$ can be defined through the following hierarchical "point-mass mixture" prior:

$$\begin{aligned} \mu_{ij} &\sim (1-\pi_{ij})\delta_0 + \pi_{ij} DP(\gamma N(0, \eta_i^2)), \\ \pi_{ij} &\sim (1-\rho_j)\delta_0 + \rho_j Beta(a, b), \\ \rho_j &\sim Beta(c, d). \end{aligned}$$

Thus in our model, not only are samples assigned to different groups (i.e., $\boldsymbol{\mu}_i = \boldsymbol{\mu}_j, 1 \le i \ne j \le n$ with positive probability), the nonzero components of the mean specific to a cluster are also clustered (i.e., $\mu_{ij} = \mu_{ik}, 1 \le j \ne k \le p$ with positive

probability). In this paper, we choose to use a more parsimonious model $\eta_i \equiv \eta$. The prior structure presented above has individual probability π_{ij} that attribute j has a nonzero effect for cluster c to which the i-th sample is assigned, while the attribute specific parameter ρ_j indicates the sparsity propensity of the covariate j. Marginalization over π_{ij} gives the more traditional point-mixture prior

$$\mu_{ij} \sim (1 - \frac{a}{a+b}\rho_j)\delta_0 + \frac{a}{a+b}\rho_j DP(\gamma N(0, \eta^2)).$$

Similar structure has been used in the regression context in Lucas et al. (2006); Seo et al. (2007); Carvalho et al. (2008). As discussed in those papers, the extended model is able to more adequately shrink towards zero through the induction of zeros for π_{ij} and thus can better separate real signals from noise and reduce false discovery.

Finally, we describe the choice of hyperpriors and the setting of hyperparameters. The base measure of the DP prior for μ_j is set as a normal distribution with $\mu_0 = y_{.j}, \sigma_0^2 = \sum_{j=1}^{p}(y_{.j} - \bar{y})^2/p$ where $y_{.j} = \sum_i y_{ij}/n$ is the observation mean for attribute j and $\bar{y} = \sum_j y_{.j}/n$ is the overall mean of all observations. For the base measure of DP prior for σ_j^2, we use the vague prior $Inv - Gamma(0.5, 0.5)$. Similarly, the standard vague conjugate prior $Inv - Gamma(0.5, 0.5)$ is also used as prior for η^2. For the four concentration parameters in the DPs, $\tau, \alpha, \beta, \gamma$, $Gamma(0.5, 0.5)$ is used as the prior. Finally, in the point-mass mixture prior, we follow Lucas et al. (2006) and set $a = 9, b = 1, c = 0.2, d = 199.8$.

At the end of this section, we emphasize that although our specification of Dirichlet process mixture model has a nested structure in that the cluster component determined by the DP mixture has a base measure on a p-dimensional vector whose component has as its distribution a mixture of zero-point mass and a Dirichlet process, the model is entirely different from the so-called nested DP (Rodriguez et al., 2008). In a nested DP, the base measure is itself a Dirichlet process. The application

in mind for nested DP is clustering of clinic centers with the goal of identifying groups of centers with similar patient outcome distributions.

3 Posterior Computation

Let $c_i, 1 \leq i \leq n$, be the latent class indicator associated with sample i, with the specific numbering of no significance, although in the presentation of the algorithm, we assume $1 \leq c_i \leq K$ when K clusters are non-empty. Similarly, c_j^μ and c_j^σ is the cluster indicator for the base-line mean and variance μ_j and σ_j^2. We also use $c_{cj}, j = 1, \ldots, p$, to indicate the clustering structure for the p components of the mean vector specific to the c-th cluster $\boldsymbol{\mu}_c = (\mu_{c1}, \ldots, \mu_{cp})$ with $c_{cj} = 0$ indicating $\mu_{cj} = 0$ is generated from the zero point mass. Similarly as before, we say $c_{ij} = c_{cj}$ if sample i is assigned to cluster c. After this augmentation of data, we can update each of the unknowns iterating between the following steps.

1. For $j = 1, \ldots, p$, draw a new value for c_j^μ using the following conditional probabilities

$$P(c_j^\mu = c|-) \propto n_{-j,c}^\mu N(\sum_{1\leq i\leq n}(y_{ij} - \mu_{ij})/n|u, 1/v + \sigma_j^2/n),$$

$$P(c_j^\mu \neq c_l^\mu, \forall l \neq j|-) \propto \alpha N(\sum_{1\leq i\leq n}(y_{ij} - \mu_{ij})/n|\mu_0, \sigma_0^2 + \sigma_j^2/n),$$

where $N(x|\mu, \sigma^2)$ denotes the normal density evaluated at x, $u = [\mu_0/\sigma_0^2 + \sum_{1\leq i\leq n, k\in C_{-j,c}^\mu}(y_{ik} - \mu_{ik})/\sigma_k^2]/v$, $v = 1/\sigma_0^2 + \sum_{1\leq i\leq n, k\in C_{-j,c}^\mu} 1/\sigma_k^2$, $C_{-j,c}^\mu$ contains all attribute indices other than j that are assigned to cluster c and $n_{-j,c}^\mu$ is the size of $C_{-j,c}^\mu$.

Then for all $c \in \{c_1^\mu, \ldots, c_p^\mu\}$, draw μ_c from $N(u, v)$ with $u = [\mu_0/\sigma_0^2 + \sum_{1\leq i\leq n, j\in C_c^\mu}(y_{ij} - \mu_{ij})/\sigma_j^2]/v$, $v = 1/\sigma_0^2 + \sum_{1\leq i\leq n, j\in C_c^\mu} 1/\sigma_j^2$, where $C_c^\mu = \{j : 1 \leq j \leq p, c_j^\mu = c\}$.

2. For $j = 1, \ldots, p$, draw a new value for c_j^σ using the following conditional probabilities

$$P(c_j^\sigma = c|-) \propto n_{-j,c}^\sigma \frac{v^u}{\Gamma(u)} \frac{\Gamma(u + n/2)}{(v + \sum_{1 \le i \le n} z_{ij}^2/2)^{u+n/2}},$$

$$P(c_j^\sigma \neq c_l^\sigma, \forall l \neq j|-) \propto \beta \frac{\beta_0^{\alpha_0}}{\Gamma(\alpha_0)} \frac{\Gamma(\alpha_0 + n/2)}{(\beta_0 + \sum_{1 \le i \le n} z_{ij}^2/2)^{\alpha_0+n/2}},$$

where $\Gamma(.)$ is the Gamma function, $z_{ij} = y_{ij} - \mu_j - \mu_{ij}$, $u = \alpha_0 + n_{-j,c}^\sigma \times n/2$, $v = \beta_0 + \sum_{1 \le i \le n, k \in C_{-j,c}^\sigma} z_{ik}^2/2$, $C_{-j,c}^\sigma$ contains all attribute indices other than j that are assigned to cluster j and $n_{-j,c}^\sigma$ is the size of $C_{-j,c}^\sigma$.

Then for all $c \in \{c_1^\sigma, \ldots, c_p^\sigma\}$, draw σ_c^2 from $Inv - Gamma(u, v)$ with $u = \alpha_0 + n_c^\sigma \times n/2$ and $v = \beta_0 + \sum_{1 \le i \le n, j \in C_c^\sigma} z_{ij}^2/2$ where $C_c^\sigma = \{j : 1 \le j \le p, c_j^\sigma = c\}$ and n_c^σ is its size.

Now suppose in the current iteration, all non-empty clusters associated with one or more samples are indicated by $\{1, \ldots, K\}$.

3. For $c \in \{1, \ldots, K\}$, $j = 1, \ldots, p$, draw π_{cj} from the conditional distribution

$$\pi_{cj}|\mu_{cj}, \rho_j \sim (1 - \rho_j)\delta_0 + \rho_j Beta(a, b + 1) \text{ if } \mu_{cj} = 0,$$

$$\pi_{cj}|\mu_{cj}, \rho_j \sim Beta(a + 1, b) \text{ if } \mu_{cj} \neq 0.$$

4. For $j = 1, \ldots, p$, draw ρ_j from the conditional distribution

$$\rho_j|\{\pi_{cj}\}_{c=1}^K \sim Beta(c + |\Pi_j|, d + K - |\Pi_j|),$$

where $\Pi_j = \{c : \pi_{cj} > 0\}$ is the set of nonzero probabilities π_{cj} associated with attribute j and $|\Pi_j|$ is the size of the set.

5. In this step, we update the clustering assignments of the samples c_i as well as

the cluster-specific mean vector $\boldsymbol{\mu}_c$. This basically makes use of Algorithm 7 in Neal (2000) which is reproduced here for completeness.

(a) For $i = 1, \ldots, n$, if c_i is not a singleton (i.e. $c_i = c_j$ for some $j \neq i$), let $c_i^* = K + 1$ be a new cluster component with $\boldsymbol{\mu}_{c_i^*}$ drawn from the prior $\mu_{c_i^* j} \sim (1 - a\rho_j/(a+b))\delta_0 + a\rho_j/(a+b)DP(\gamma N(0, \eta^2))$. Accept c_i^* with probability

$$\min[1, \frac{\tau}{n-1}\frac{F(\mathbf{y}_i; \boldsymbol{\mu}_{c_i^*})}{F(\mathbf{y}_i; \boldsymbol{\mu}_{c_i})}],$$

where $F(\mathbf{y}_i; \boldsymbol{\mu}_c) = \prod_{j=1}^{p} N(y_{ij}|\mu_{cj}, \sigma_j^2)$.

If c_i is a singleton, propose $c_i^* = c, c \in \{c_1, \ldots, c_n\}$ with probability $n_{-i,c}/(n-1)$ ($n_{-i,c}$ is the number of samples excluding the i-th sample that are currently assigned to cluster c) and accept with probability

$$\min[1, \frac{n-1}{\tau}\frac{F(\mathbf{y}_i; \boldsymbol{\mu}_{c_i^*})}{F(\mathbf{y}_i; \boldsymbol{\mu}_{c_i})}].$$

(b) For $i = 1, \ldots, n$, if c_i is not a singleton, set $c_i = c, c \in \{c_1, \ldots, c_n\}$ with probability

$$P(c_i = c|-) \propto \frac{n_{-i,c}}{n-1}F(\mathbf{y}_i; \boldsymbol{\mu}_c).$$

(c) For all $c \in \{c_1, \ldots, c_n\}$, perform a Gibbs sampling step for the components of $\boldsymbol{\mu}_c|\{\mathbf{y}_i\}, \{\rho_j\}$ with $i \in \{1 \leq k \leq n : c_k = c\}$:

i. For $j = 1, \ldots, p$, update c_{cj} from the following distribution

$$P(c_{cj} = 0|-) \quad \propto \quad (1 - \frac{a}{a+b}\rho_j)N(x_j|0, \sigma_j^2/n_c),$$

for $c' \in \{c_{ck}, 1 \leq k \leq p, k \neq j\}$,

$$\begin{aligned} P(c_{cj} = c'|) &\propto \frac{a}{a+b}\rho_j n_{-j,c'} N(x_j|u, 1/v + \sigma_j^2/n_c), \\ P(c_{cj} \neq c_{ck} \forall k \neq j|) &\propto \frac{a}{a+b}\rho_j \gamma N(x_j|0, \eta^2 + \sigma_j^2/n_c), \end{aligned}$$

where $x_j = \sum_{i:c_i=c}(y_{ij} - \mu_j)/n_c$, $u = \sum_{j:c_{cj}=c'}(x_j/\sigma_j^2)/v$ and $v = 1/\eta^2 + \sum_{j:c_{cj}=c'} 1/\sigma_j^2$.

ii. For $c' \in \{c_{c1}, \ldots, c_{cp}\}$, draw a new value for $\mu_{cc'}$ from $\mu_{cc'}|\{y_{ij}\}$, with $c_i = c, c_{cj} = c'$.

Note that here we used a partially collapsed Gibbs step (van Dyk and Park, 2008) by integrating out π_{ij}.

6. Draw η^2 from the conditional distribution

$$\eta^2 \sim Inv - Gamma((1 + n^\mu)/2, (1 + \sum_{(c,j)\in C^\mu} \mu_{cj}^2)/2),$$

where C^μ is the set of indices of all nonzero unique values of μ_{cj}, and n^μ is the size of the set.

7. Draw the concentration parameters $\alpha, \beta, \gamma, \tau$ using the data augmentation approach of Escobar and West (1995).

A note is in order. All the updates above are obtained by standard calculations. In particular, step 5 is a reproduction of Algorithm 7 in Neal (2000), with the only difference being step 5(c) where Gibbs sampling must be used since the conditional distribution $\boldsymbol{\mu}_c$ conditional on samples assigned to cluster c is not directly available in closed form.

In step 5(a) above, if c_i is a singleton, we draw a new value for $\boldsymbol{\mu}_c$ from the prior, which is a high-dimensional vector, and the prior distribution is a mixture of zero point mass and a Dirichlet process. This typically makes $F(\mathbf{y}_i; \boldsymbol{\mu}_c)$ extremely small, which is not surprising since $\boldsymbol{\mu}_c$ drawn from the prior can hardly explain the observed sample $\mathbf{y}_i$ well. When the update is implemented as presented, this proposal is almost never accepted. To solve this problem, we successfully used a sequential sampling approach that proposes a new value for $\boldsymbol{\mu}_c$ taking into account the observed $\mathbf{y}_i$. The proposed sequential sampling step generates the new value with the following scheme:

- With $c = K + 1$, for $j = 1, \ldots, p$, draw c_{cj} from the following distribution

$$\begin{aligned} P(c_{cj} = 0|-) &\propto (1 - \frac{a}{a+b}\rho_j) N(x_j|0, \sigma_j^2/n_c), \\ P(c_{cj} = c'|-) &\propto \frac{a}{a+b}\rho_j n_{-j,c} N(x_j|u, 1/v + \sigma_j^2/n_c),\ c' \in \{c_{ck} : k \neq j\}, \\ P(c_{cj} \neq c_{ck} \forall k \neq j|-) &\propto \frac{a}{a+b}\rho_j \alpha N(x_j|0, \eta^2 + \sigma_j^2/n_c), \end{aligned}$$

  where we conditioned on $\{\mu_{ck}\}_{k=1}^{j-1}, x_j, \{\sigma_k\}_{k=1}^{p}, \rho_j$ and η. In the above we define $x_j = \sum_{i:c_i=c}(y_{ij} - \mu_j)/n_c$, $u = \sum_{j:c_{cj}=c'} x_j/\sigma_j^2/v$ and $v = 1/\eta^2 + \sum_{j:c_{cj}=c'} 1/\sigma_j^2$.

- For $c' \in \{c_{k1}, \ldots, c_{kp}\}$, draw a new value for $\mu_{cc'}$.

These expressions are very similar to the Gibbs step 5(c), with the important difference that when proposing new value for c_{cj}, only the previously sampled $c_{ck}, k < j$ are available and the update is performed sequentially. With this change of proposal distribution, the acceptance probability in step 5(a) should be changed to

$$\min[1, \frac{\tau}{n-1} \frac{F(\mathbf{y}_i; \boldsymbol{\mu}_{c_i^*})}{F(\mathbf{y}_i; \boldsymbol{\mu}_{c_i})} \frac{Q_0(\boldsymbol{\mu}_{c_i^*})}{Q(\boldsymbol{\mu}_{c_i^*})}]$$

and

$$\min[1, \frac{n-1}{\tau}\frac{F(\mathbf{y}_i;\boldsymbol{\mu}_{c_i^*})}{F(\mathbf{y}_i;\boldsymbol{\mu}_{c_i})}\frac{Q(\boldsymbol{\mu}_{c_i})}{Q_0(\boldsymbol{\mu}_{c_i})}]$$

respectively, where Q_0 is the proposal density when $\boldsymbol{\mu}_c$ is drawn from the prior and Q is the proposal density of $\boldsymbol{\mu}_c$ when it is drawn from the sequential sampling approach described above.

Finally, for inferences on cluster-specific parameters, such as μ_{cj} and π_{cj}, we need to take care of the label switching problem (Stephens, 2000; Tadesse et al., 2005). This can be done conditionally on the number of clusters for the samples, K. We refer the readers to Tadesse et al. (2005) for details.

4 Extension to Variance Shifts

Although it is not our focus in this paper, our model can easily be extended to the case where groups are distinguished by both different mean and variance for one or more attributes. On can extend the model and write the observed data as

$$y_{ij} = \mu_j + \mu_{ij} + \sigma_j\sigma_{ij}\epsilon_{ij},\ \epsilon_{ij} \overset{i.i.d.}{\sim} N(0,1).$$

The extra factor σ_{ij} indicates the difference in variance for data assigned to different groups. Similarly to the mean shift model, we can put a hierarchical sparsity prior

$$\begin{aligned}
\sigma_{ij}^2 &\sim (1-\pi_{ij}^\sigma)\delta_1 + \pi_{ij}^\sigma DP(\kappa Inv-Gamma(\alpha_\sigma,\beta_\sigma)),\\
\pi_{ij}^\sigma &\sim (1-\rho_j^\sigma)\delta_0 + \rho_j^\sigma Beta(a,b),\\
\rho_j^\sigma &\sim Beta(c,d).
\end{aligned}$$

Note that for variances, sparsity means many of the σ_{ij} will be exactly equal to one. For the base measure, we should also choose it to be roughly centered at one for

identifiability, for example, we can use $\alpha_\sigma = 1.5, \beta_\sigma = 0.5$ as a vague prior.

This extension causes few extra complications on the updating strategy for posterior computation, with extra updates for σ_{ij} as well as some slight change of formula. The details are omitted here. We do not consider further the case with variance shift since one can argue that for the microarray analysis for example, the researchers typically focus on mean shift as a distinguishing feature of tissue subtypes.

5 Simulation and Application

5.1 Simulation Study

We investigate the performance of our estimation method in a simulation study. A dataset containing 20 samples and 200 covariates is generated as follows.

$$
\begin{aligned}
y_{ij} &= \mu_{ij} + \sigma_j \epsilon_{ij},\ \epsilon_{ij} \sim N(0,1), \\
\mu_{i,j} &= 0.25, 1 \le i \le 5, 1 \le j \le 5, \\
\mu_{i,j} &= 0.1, 6 \le i \le 10, 1 \le j \le 5, \\
\mu_{i,j} &= -0.1, 11 \le i \le 15, 1 \le j \le 5, \\
\mu_{i,j} &= -0.25, 16 \le i \le 20, 1 \le j \le 5, \\
\mu_{i,j} &= 0.2, 1 \le i \le 5, 6 \le j \le 10, \\
\mu_{i,j} &= -0.15, 16 \le i \le 20, 11 \le j \le 15, \\
\mu_{ij} &= 0 \text{ otherwise}, \\
\sigma_j &= 0.1, 1 \le j \le 15, \\
\sigma_j &= 0.05, \text{ otherwise}.
\end{aligned}
$$

The structure of μ_{ij} is shown in Figure 1(a) where different values for μ_{ij} show up as different gray levels. Each row in the image represents a sample and each column represents an attribute. Only the first 50 attributes are shown. The first 5 covariates distinguishes across all four groups, while attributes 6-10 distinguish the first cluster from the others and attributes 11-15 distinguish the fourth cluster from the others. We use the model described in Section 2 to fit the simulated dataset. Figure 1(b) shows the observed data in the same format as Figure 1(a). Figure 2 shows the posterior updates for the number of clusters identified as well as gives some indication of the mixing of the Markov chain. The posterior gives strong support for four clusters, with support for five clusters comes next. In simulation as well as real data application that follows, we used a burn-in period of $10,000$ updates and $40,000$ iterations after burn-in for inferences. The posterior estimates of μ_{ij} is shown in Figure 1(c) as a matrix for the first 50 attributes only. Four clusters and the zero structures are clearly identified. In contrast, the hierarchical clustering based on COSA algorithm of Friedman and Meulman (2004) failed to identify the true clusters, as shown in Figure 3 for single, average, and complete linkage.

In Figure 4, we show the posterior estimate of ρ_j for $1 \leq j \leq 50$, which indicates the contribution of the j-th attribute to cluster discrimination. The results are quite encouraging, with the first 15 attributes clearly identified as signal variables and the first 5 attributes estimated to be associated with larger values of ρ_j, consistent with the simulation scheme. We can also use a simple threshold of 0.5 on posterior estimates of π_{ij}. In particular, we decide attribute j to be relevant for clustering if $\pi_{ij} > 0.5$ for at least one i. This strategy also exactly identifies the first 15 attributes as significant.

Finally, for this simulated example, using DP for μ_j and $\mu_{ij}, 1 \leq j \leq p$ performs better than a normal prior (corresponding to the case with $\alpha \to \infty$ and $\gamma \to \infty$. The mean squared error of $\mu_j + \mu_{ij}, 1 \leq i \leq 20, 1 \leq j \leq 15$, under our model is 0.006, in

contrast with 0.011 when $\alpha, \gamma \to \infty$. This is consistent with the results reported in Nott (2008).

Now we use additional simulated examples with various choices of the number of attributes p and the values for mean vectors to investigate the performance of our method. Besides the example above, we use the following simulation schemes. All examples are simulated from model $y_{ij} = \mu_{ij} + \sigma_j \epsilon_{ij}$ as before.

- Example 2. Same as the example presented above, except the number of noise variables are increased to $p = 1000$.

- Example 3. Here we have $n = 20$ samples and $p = 50$ attributes, among which 10 attributes are informative for clustering across all four clusters. The cluster sizes are chosen to be $3, 3, 7, 7$ respectively, with $\mu_{ij} = c/4$ if sample i belongs to cluster c, $1 \leq c \leq 4$, when $1 \leq j \leq 10$, and $\sigma_j = 0.1$ for all j.

- Example 4. Here $n = 20$ and $p = 50$, with $\mu_{ij} = j/50$ when $1 \leq i \leq 10$, $\mu_{ij} = (50 - j)/50$ when $11 \leq i \leq 20$, and $\sigma_j = 0.1$ for all j.

Our method is compared to two other model-based Bayesian clustering methods with variable selection proposed in Hoff (2006) and Kim et al. (2006). For each example, we simulated 50 datasets. The different methods are compared using three performance measures. First, we compute the average mean squared errors for μ_{ij} where we only consider attributes that are relevant for clustering. Also, we consider the number of attributes selected by each method as well as its overlap with true relevant attributes. Based on the existing implementation for Hoff's approach, the attributes are selected such that they maximize the joint posterior distribution. Besides, all three approaches can identify the correct number of clusters in all examples.

The results in Table 1 show that our approach always outperformed the other two for all the examples considered here in terms of mean squared error. Examples

1 through 3 are perhaps favorable to our approach. Example 4 was intended for a situation where all components for a cluster mean are distinct. Our method still works well in this situation in terms of mean squared error, compared with the other two approaches. In terms of selected variables, the approach of Hoff (2006) tends to select a large number of relevant variables resulting in high false discovery rate. When there exist attributes that only distinguish a subset of cluster components, such as the situation in Examples 1 and 2, the variable selection approach of Kim et al. (2006) tends to miss some of the those covariates. Note that the approach of Kim et al. (2006) cannot give informations on whether an attribute only distinguishes a subset of the cluster components even if the attribute is selected.

5.2 Leukemia Gene Expression Data Example

We use the leukemia gene expression dataset (Golub et al., 1999) to demonstrate the utility of our proposed method. The training dataset contains 38 tissue samples, among which 11 samples are acute myeloid leukemia (AML) and the rest are acute lymphoblastic leukemia (ALL). The 27 ALL samples are further divided into two subgroups: 8 T-cell and 19 B-cell samples. The samples were arrayed with a total of 7129 genes in a microarray experiment. Following the standard preprocessing steps in Dudoit et al. (2002), we truncate the expression values to within the interval $[1, 16000]$, and delete those genes whose maximum and minimum expression across all samples satisfies $\max / \min \leq 5$ and $\max - \min \leq 500$. Finally, we select the top 2000 genes with the largest variances across all samples so that at the end we have for this dataset $n = 38, p = 2000$.

We apply our proposed method to the dataset with the hyperparameters set exactly as discussed in Section 2. Convergence of the MCMC updates is invariably a concern in high-dimensional problem with variable selection. As a simple diagnostic,

two MCMC runs of 50,000 iterations with the first 10,000 as burn-in are implemented, with different initialization. In particular, we start one Markov chain with initially all samples assigned to one cluster, and another chain where each sample is assigned to its own separate cluster. The posterior estimates of various unknown quantities for the two runs shows good agreement which indicates the chains mixed well in our implementation.

As shown in Figure 5, the posterior for this dataset put most of the support for the number of clusters between 3 and 9, with 6 clusters receiving the highest score. Conditional on $K = 6$, setting the threshold 0.5 for the posterior estimates of π_{cj} returns 872 genes. This is much larger than the 120 genes reported in Kim et al. (2006). Previous studies, such as Thomas et al. (2001), also demonstrated that there were a large number of genes differentially expressed between different tissue samples.

For inference about the cluster structure and comparison to known tissue subtypes, we estimate the posterior probability of $c_i = c, 1 \leq c \leq 6$ from posterior samples conditioned on $K = 6$ with the help of the procedure that deals with label switching. Each sample is allocated to the cluster with the largest posterior probability. The relationship between this allocation and known tissue types are shown in Table 2. Under our method, one of the ALL samples is misclassified into a cluster dominated by AML samples, also one of the AML samples is misclassified into a cluster dominated by ALL samples. We see that the known AML and ALL-B tissue types might further consist of some subtypes. Using our method, we can also discover genes that distinguish only some subgroups. For example, among those 872 genes relevant for clustering only 64 of them can distinguish between ALL and AML samples without discriminative power for different subtypes.

6 Conclusion

In this article, we propose a novel Bayesian approach to high-dimensional clustering with variable selection. The distinguishing features of our method include a separate Dirichlet process for shrinkage estimation of cluster mean, as well as a hierarchical point-mass structure that improves the separation of significant signal from noise variables. We propose a sequential sampling approach in one of the updating iterations of the MCMC algorithm to solve the computational problem associated with the high dimensionality of the mean vector.

Our approach only involves diagonal covariance matrices. It has been argued in other studies in both supervised and unsupervised context that for "high dimensional low sample size" setting, this working independence assumption is more effective than the full covariance matrix approach (Fraley and Raftery, 2006; Bickel and Levina, 2004; Tibshirani et al., 2003). Generalization of our method to general covariance structure seems much more challenging.

In our implementation, we choose to use Algorithm 7 in Neal (2000) for its simplicity in implementation. More efficient approaches like split-merge update (Jain and Neal, 2004) can also be utilized. Nevertheless, we still expect the original algorithm should be modified using sequential sampling instead of drawing new components from the prior to make the implementation feasible in practice.

References

Banfield, J. D. and Raftery, A. E. "Model-based Gaussian and non-Gaussian clustering." *Biometrics*, 49(3):803–821 (1993).

Bickel, P. J. and Levina, E. "Some theory for Fisher's linear discriminant function,

'naive Bayes', and some alternatives when there are many more variables than observations." *Bernoulli*, 10(6):989–1010 (2004).

Carvalho, C., Chang, J., Lucas, J., Nevins, J., Wang, Q., and West, M. "High-dimensional sparse factor modeling: applications in gene expression genomics." *Journal of the American Statistical Association*, 103(484):1438–1456 (2008).

Dudoit, S., Fridlyand, J., and Speed, T. P. "Comparison of discrimination methods for the classification of tumors using gene expression data." *Journal of the American Statistical Association*, 97(457):77–87 (2002).

Escobar, M. D. and West, M. "Bayesian density estimation and inference using mixtures." *Journal of the American Statistical Association*, 90(430):577–588 (1995).

Fraley, C. and Raftery, A. E. "Model-based clustering, discriminant analysis, and density estimation." *Journal of the American Statistical Association*, 97(458):611–631 (2002).

—. "MCLUST version 3 for R: Normal mixture modeling and model-based clustering." *Technical Report no. 504, Department of Statistics, University of Washington* (2006).

Friedman, J. H. and Meulman, J. J. "Clustering objects on subsets of attributes." *Journal of the Royal Statistical Society Series B-Statistical Methodology*, 66:815–839 (2004).

George, E. I. and McCulloch, R. E. "Variable selection via Gibbs sampling." *Journal of the American Statistical Association*, 88(423):881–889 (1993).

—. "Approaches for Bayesian variable selection." *Statistica Sinica*, 7(2):339–373 (1997).

Golub, T. R., Slonim, D. K., Tamayo, P., Huard, C., Gaasenbeek, M., Mesirov, J. P., Coller, H., Loh, M. L., Downing, J. R., Caligiuri, M. A., Bloomfield, C. D., and Lander, E. S. "Molecular classification of cancer: Class discovery and class prediction by gene expression monitoring." *Science*, 286(5439):531–537 (1999).

Hoff, P. "Model-based subspace clustering." *Bayesian analysis*, 1:321–344 (2006).

Jain, S. and Neal, R. M. "A split-merge Markov chain Monte Carlo procedure for the dirichlet process mixture model." *Journal of Computational and Graphical Statistics*, 13(1):158–182 (2004).

Kim, S., Tadesse, M. G., and Vannucci, M. "Variable selection in clustering via Dirichlet process mixture models." *Biometrika*, 93(4):877–893 (2006).

Liu, J., Zhang, J., Palumbo, M., and Lawrence, C. "Bayesian clustering with variable and transformation selection." *Bayesian Statistics*, 7:249–275 (2003).

Lucas, J., Carvalho, C., Wang, Q., Bild, A., Nevins, J. R., and West, M. "Sparse statistical modelling in gene expression genomics." In *Bayesian Inference for Gene Expression and Proteomics*, 155–176. Cambridge University Press (2006).

Neal, R. M. "Markov chain sampling methods for Dirichlet process mixture." *Journal of Computational and Graphical Statistics*, 9(2):249–265 (2000).

Nott, D. J. "Predictive performance of Dirichlet process shrinkage methods in linear regression." *Computational Statistics & Data Analysis*, 52(7):3658–3669 (2008).

Rodriguez, A., Dunson, D. B., and Gelfand, A. E. "The nested Dirichlet process." *Journal of the American Statistical Association*, 103(483):1131–1144 (2008).

Seo, D. M., Goldschmidt-Clermont, P. J., and West, M. "Of mice and men: Sparse statistical modeling in cardiovascular genomics." *Annals of Applied Statistics*, 1(1):152–178 (2007).

Stephens, M. "Dealing with label switching in mixture models." *Journal of the Royal Statistical Society Series B-Statistical Methodology*, 62:795–809 (2000).

Tadesse, M. G., Sha, N., and Vannucci, M. "Bayesian variable selection in clustering high-dimensional data." *Journal of the American Statistical Association*, 100(470):602–617 (2005).

Thomas, J. G., Olson, J. M., Tapscott, S. J., and Zhao, L. P. "An efficient and robust statistical modeling approach to discover differentially expressed genes using genomic expression profiles." *Genome Research*, 11(7):1227–1236 (2001).

Tibshirani, R. "Regression shrinkage and selection via the Lasso." *Journal of the Royal Statistical Society Series B-Methodological*, 58(1):267–288 (1996).

Tibshirani, R., Hastie, T., Narasimhan, B., and Chu, G. "Class prediction by nearest shrunken centroids, with applications to DNA microarrays." *Statistical Science*, 18(1):104–117 (2003).

van Dyk, D. A. and Park, T. "Partially collapsed Gibbs samplers: Theory and methods." *Journal of the American Statistical Association*, 103(482):790–796 (2008).

Table 1: Simulation results for the four examples. The numbers in each cell correspond to median mean squared error, number of attributes selected and overlap between selected attributes and the truth, respectively.

Methods	Example 1	Example 2	Example 3	Example 4
Our approach	0.005 15 15	0.021 18 13	0.012 8 8	0.016 46 46
Hoff (2006)	0.013 19 15	0.046 26 10	0.027 17 7	0.025 50 49
Kim et al. (2006)	0.018 14 14	0.069 14 8	0.016 11 10	0.021 48 48

Table 2: Clustering results for leukemia expression data conditional on $K = 6$.

	cluster from the proposed method					
samples	1	2	3	4	5	6
ALL-T(8)	0	0	8	0	0	0
ALL-B(19)	0	1	0	6	4	8
AML(11)	7	3	0	0	1	0

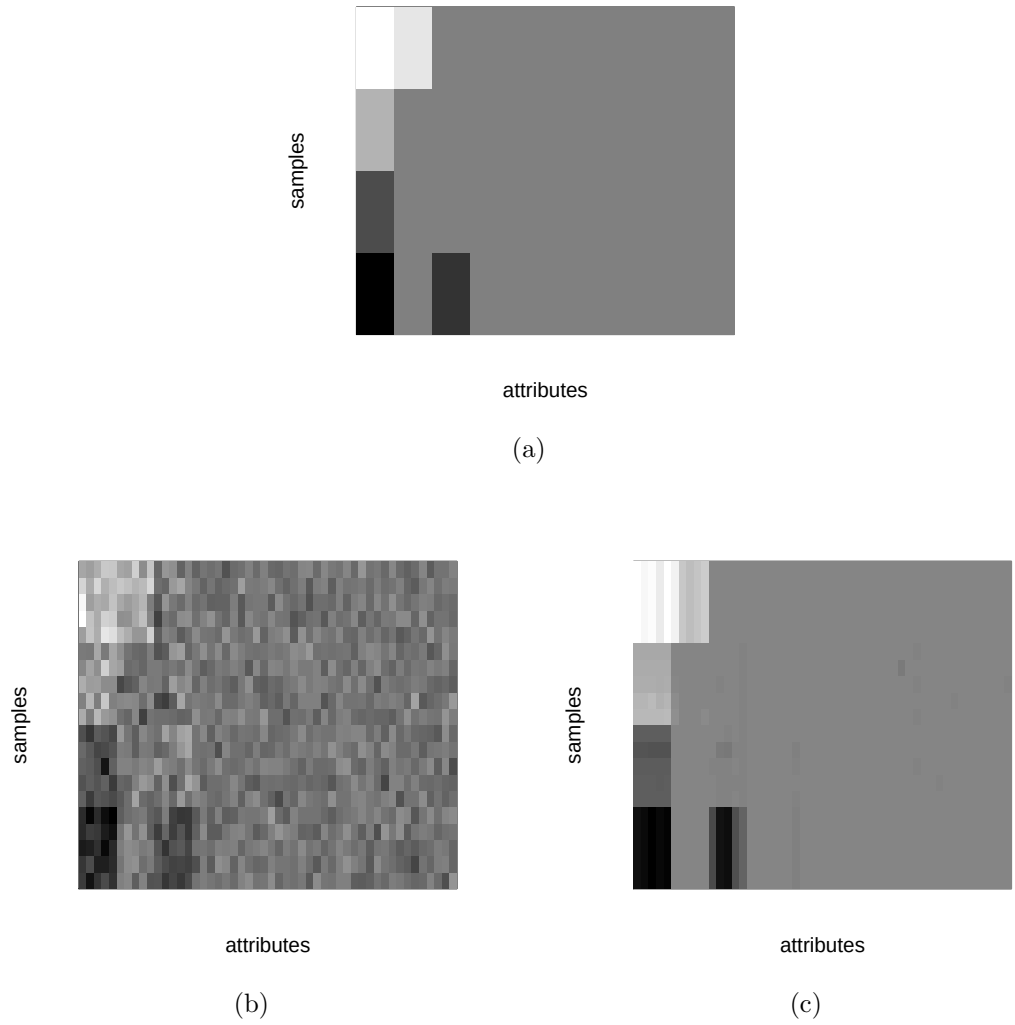

Figure 1: (a) Mean structure μ_{ij} for the simulated data plotted as an image. (b) Noisy observed data. (c) Estimated mean under our approach.

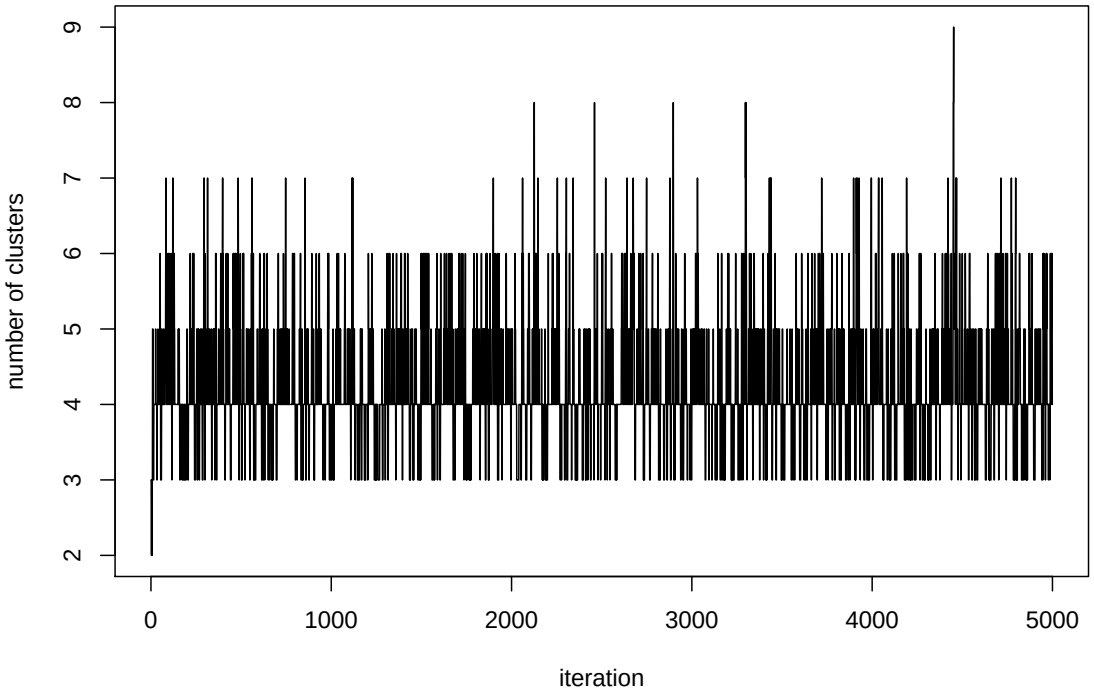

Figure 2: Trace plot of the number of clusters K for the simulated dataset.

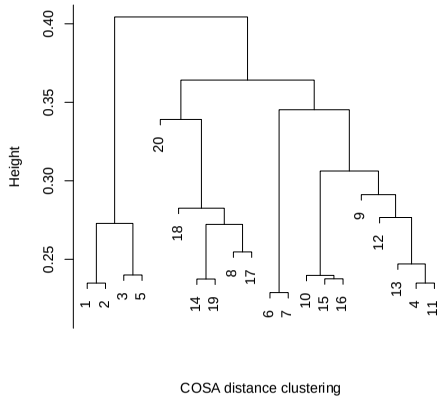

(a)

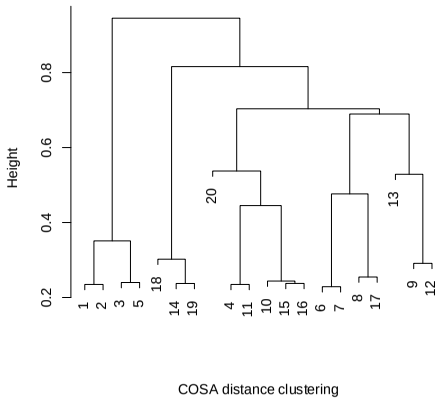

(b)

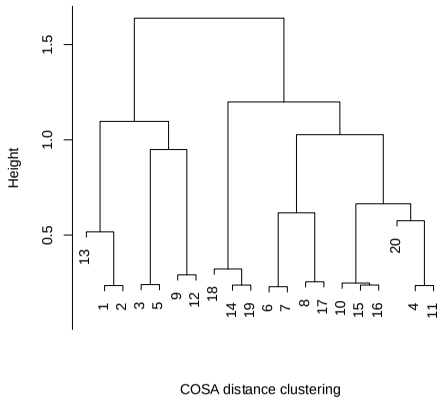

(c)

Figure 3: Results from hierarchical clustering with COSA algorithm based on (a) single linkage; (b) average linkage; (c) complete linkage.

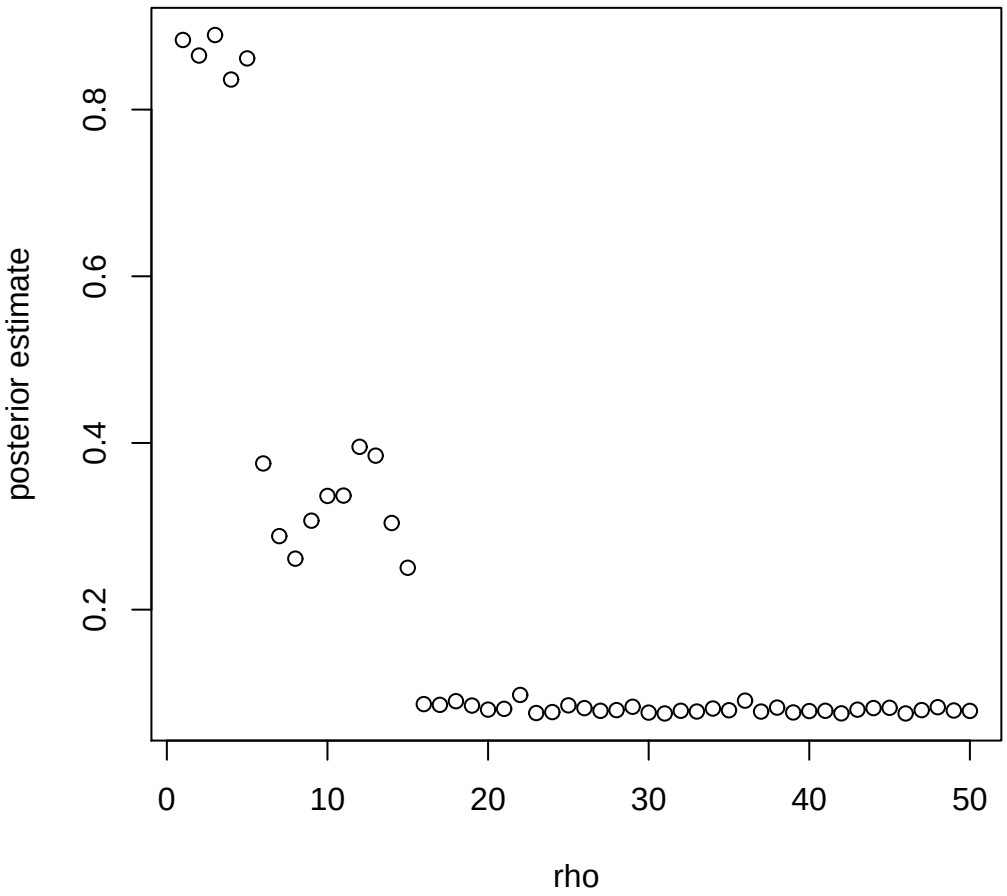

Figure 4: Posterior estimates of ρ_j for $1 \leq j \leq 50$.

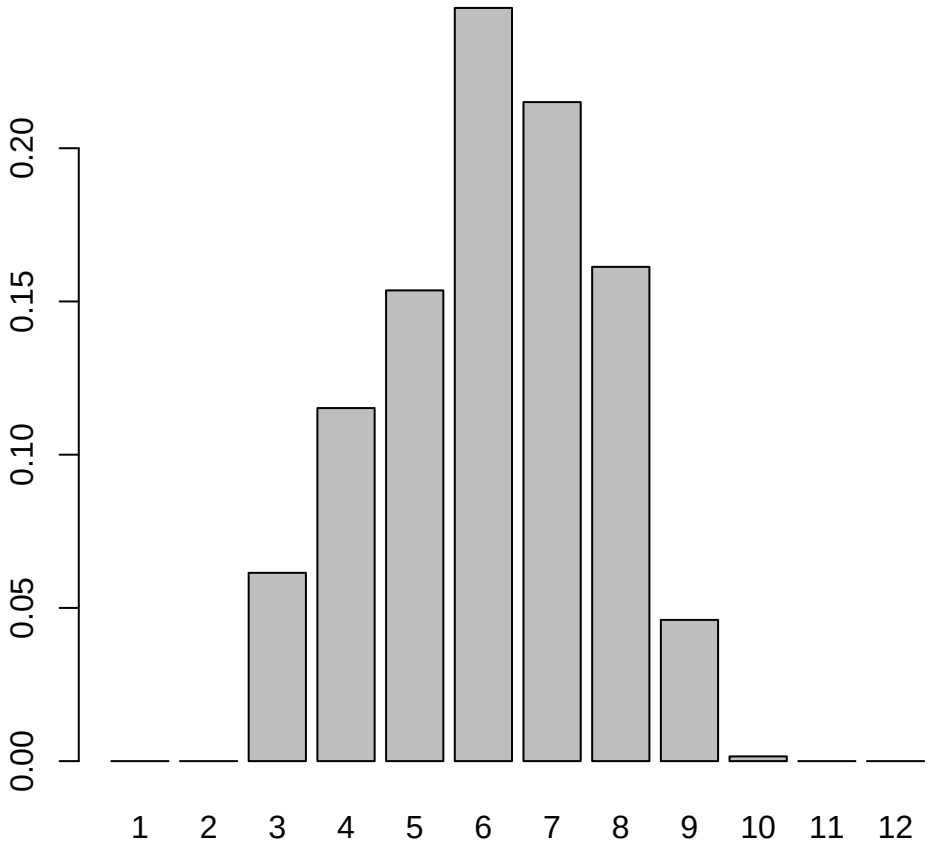

Figure 5: Posterior estimate of the number of clusters for the leukemia expression dataset.